

2014-01-01

Multi-objective routing optimisation for battery-powered wireless sensor mesh networks

Rahat, AAM

<http://hdl.handle.net/10026.1/12585>

10.1145/2576768.2598311

GECCO 2014 - Proceedings of the 2014 Genetic and Evolutionary Computation Conference

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Multi-Objective Routing Optimisation for Battery-Powered Wireless Sensor Mesh Networks

Alma A. M. Rahat
Computer Science
University of Exeter, UK
A.A.M.Rahat@exeter.ac.uk

Richard M. Everson^{*}
Computer Science
University of Exeter, UK
R.M.Everson@exeter.ac.uk

Jonathan E. Fieldsend
Computer Science
University of Exeter, UK
J.E.Fieldsend@exeter.ac.uk

ABSTRACT

Mesh network topologies are becoming increasingly popular in battery powered wireless sensor networks, primarily due to the extension of network range and resilience against routing failures. However, multi-hop mesh networks suffer from higher energy costs, and the routing strategy directly affects the lifetime of nodes with limited energy sources. Hence while planning routes there are trade-offs to be considered between individual and system-wide battery lifetimes. We present a novel multi-objective routing optimisation approach using evolutionary algorithms to approximate the optimal trade-off between minimum lifetime and the average lifetime of nodes in the network. In order to accomplish this combinatorial optimisation rapidly and thus permit dynamic optimisation for self-healing networks, our approach uses novel k -shortest paths based search space pruning in conjunction with a new edge metric, which associates the energy cost at a pair of nodes with the link between them. We demonstrate our solution on a real network, deployed in the Victoria & Albert Museum, London. We show that this approach provides better trade-off solutions in comparison to the minimum energy option, and how a combination of solutions over the lifetime of the network can enhance the overall minimum lifetime.

Categories and Subject Descriptors

C.2.1 [Computer Systems Organization]: Communication networks—*network architecture and design, wireless communication*; I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search—*Graph and tree search strategies, Heuristics methods*; G.2.2 [Discrete Mathematics]: Graph Theory—*Network problems, Path and circuit problems*

General Terms

Algorithms, Experimentation, Performance, Reliability

^{*}Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.
Copyright 2014 ACM 978-1-4503-2662-9/14/07 ...\$15.00.
<http://dx.doi.org/10.1145/2576768.2598311>.

Keywords

Mesh networks, shortest path, evolutionary algorithms, multi-objective optimisation, network lifetime optimisation

1. INTRODUCTION

Wireless Sensor Networks (WSNs) consist of autonomous devices distributed over a wide area that are able to sense and periodically report environmental parameters such as temperature or humidity. They are extensively used for remote monitoring, especially in industrial and regulatory applications. Many applications require sensors to be placed far from easy access to mains power so battery-powered sensors are popular and are often necessary in such situations. However, it is also desirable that sensors can be left unattended, without battery replacement for as long as possible. In this paper we therefore investigate the use of multi-objective evolutionary algorithms to find routing schemes for *mesh* wireless sensor networks that optimally preserve the life of the network.

Many existing commercial products use an *ad hoc* topology for WSNs. Generally, these *ad hoc* networks are point to multi-point networks, in which each sensor reports directly to a central base station. Despite being a low power solution, major limitations include limited network range as each sensor node must be able to communicate directly with the base station, and the strictly rigid network structure means that networks are unable to cope with changing conditions in a dynamic radio environment, which may occur as buildings are modified or furniture is moved. As a consequence sometimes multiple base stations are required to cover a single site.

As a solution, in recent years, mesh network topologies, where data is relayed from node to node *en route* to the base station, have grown in popularity. In addition to extending the range that can be covered by using multiple hops, mesh networks provide the opportunity of using alternative routes, thus increasing resilience to radio environment changes. However, mesh networks are expensive in terms of energy consumption due to a higher overhead at each node for additional activities, namely relaying messages for other nodes and, in systems with distributed planning, calculating new routes. These additional activities can be severely detrimental to the overall life of the network before it requires servicing and battery replacement. It is desirable to maximise the average life of nodes in the network, but in addition it will usually be important to maximise the time before the battery of the shortest-lived node is exhausted. Therefore, routing optimisation in mesh networks with careful consid-

eration of the trade-off between individual and system-wide battery lifetimes within the network, is necessary for a practical implementation.

Efficient energy usage in WSNs has attracted much research interest in recent years, especially because frequent changing of batteries is infeasible. Approaches can be divided into distributed and centralised [10]. Distributed approaches, in which responsibility for routing is distributed across the constituent nodes, involve techniques like reinforcement learning [7] and swarm intelligence [1], etc. On the other hand centralised systems, in which a centrally calculated route is broadcast to participating nodes, mostly incorporate variants of evolutionary algorithms (EAs) [11, 13, 8]. In the distributed approach, nodes need to have sufficient computational power to decide on the best path for sending data and sufficient storage for storing information about local connectivity. In comparison centralised approaches require lower computational power and lower static storage for storing route information because they only require knowledge about the node specific route information, but there is system-wide overhead incurred for gathering connectivity information and broadcasting routing information. In this paper we consider very low powered nodes each of which has limited computational power and storage. Routes are therefore computed at a mains-powered base station node.

Most current EA-based energy-aware centralised systems consider energy expense. This is the case even for multi-objective routing optimisation, where energy expense is optimised with other objectives based on factors like quality of service, bandwidth, packet loss ratio, etc. [11, 13]. However, optimising the overall energy expenditure of a network may be detrimental to the overall performance of the network, because often the goal is to prolong the lifetime of the network before batteries need replacing. Merely reducing the overall energy expenditure may place a large burden on a few nodes, resulting in the rapid exhaustion of their batteries. We therefore seek to optimise the lifetime of network nodes by modelling the charge held in their batteries and the energy expenditure at *each* node. Islam *et al.* [8], and Kamath *et al.* [9] have considered maximising only the minimum remaining lifetime among nodes. Such approaches can improve individual node specific energy state, but can be sub-optimal from system-wide perspective. We therefore seek to find the optimal trade-off between local and network-wide battery lifetimes.

In this paper, we consider a centralised approach due to the limitation of computational power and storage that a real network entails. As network longevity is our primary goal, we use a multi-objective evolutionary optimiser to simultaneously maximise the minimum lifetime of any node in the network together with the average lifetime of all nodes in the network. In practice our system also monitors packet-loss-ratio and quality-of-service performance indicators, so that re-optimisation can be invoked if performance declines, and thus permitting self-healing in dynamic radio environments. Using this procedure we are able to achieve better performance than minimum energy route and we show that using a combination of solutions over the lifetime of the network can extend minimum lifetime of the network.

In section 2 we describe the model of the system. In section 3 we present a novel search space pruning method. In section 4 we discuss the multi-objective evolutionary routing optimisation strategy. Based on this, in section 5, we




Figure 1: Route $S_i = \langle v_i, v_j, \dots, v_B \rangle$ from v_i via v_j to the base station v_B .

demonstrate and discuss our findings in a real network deployed in the Victoria & Albert Museum, London. In section 6 we discuss how minimum lifetime of the system can be extended through re-optimisation. In section 7 we present re-optimisation as a tool to cope with the dynamism of the radio environment. Finally, the conclusions are presented in section 8.

2. SYSTEM MODEL AND PARETO OPTIMALITY

In this section we model the WSN and derive a multi-objective problem in order to investigate the trade-off between individual and system-wide impacts on the remaining lifetimes of nodes due to possible routing schemes in a network.

A WSN is represented as a network graph, $G = \{V, E\}$, where V is a finite set of n sensor nodes v_i plus a base station node, v_B , and E is the finite set of m edges [2, 4]. Each node must communicate with the base station, perhaps by relaying a message through one or more other nodes. In our scheme each node reports its status once every reporting cycle (e.g. once each minute).

As illustrated in Figure 1, a route from node v_i to the base station v_B is described by the sequence, $S_i = \langle v_i, v_j, \dots, v_B \rangle$. We denote by $S_i[p]$ the p th element of the route S_i . A *routing scheme* R is a set of routes, one for each node in the network to the base station:

$$R = \{S_1, S_2, \dots, S_n\}. \quad (1)$$

An initial mapping phase, preceding optimisation, is used to discover which other nodes the node v_i can communicate with. Communications may take place using a variety of baud rates and powers, so we assume that the most energy efficient baud rate and power combination has been discovered for each pair of nodes that can communicate; generally low power and high baud rates are most efficient in our system, but the optimisation does not rely on this. Also we assume that communication is reliable at this baud rate and power combination (given the current configuration of the physical environment in which the network is deployed).

The energy required to transmit a message from v_i to the base station is found by summing the energies required to transmit a message between each of the nodes comprising the route:

$$H_i = \sum_{p=1}^{l-1} e_{S_i[p], S_i[p+1]} \quad (2)$$

where l is the length of the route and e_{jk} is the energy required to transmit a message from v_j to v_k . Note that this generally involves energy expenditure at both the transmitting node and the receiving node, and will also involve expenditures for transmitting an acknowledgement. As noted above, we assume that the communication is reliable, but if an acknowledgement is not received from the receiver the message is resent; this additional expense is not modelled, but if a link becomes unreliable, the routing is re-optimised.

In many routing optimisation problems, such as shortest path problems, minimising a route's overall cost is desirable. The overall cost is found by summing the costs associated with each edge in the route. There are many well-known methods for minimising such costs, e.g. [5]. In this problem, however, we focus on the costs expended at the *nodes* themselves, rather than the edge costs. This is because it is energy expended at the nodes that depletes charge in the batteries and thus governs the lifetime of a node.

Let T_{ij} be the energy (charge) required at node v_i to send a message to v_j and let A_{ik} be the energy required to receive a message from v_k at v_i . Then in one reporting cycle, the energy expended at v_i in sending its own data to its downstream node $d = S_i$ [2] and relaying messages received from nodes with indices in the set \mathcal{I} and sending them on to nodes with indices \mathcal{O} is

$$C_i = T_{id} + \sum_{k \in \mathcal{I}} A_{ik} + \sum_{j \in \mathcal{O}} T_{ij}. \quad (3)$$

Clearly $\sum_i^n C_i = \sum_i^n H_i$ is equal to the energy cost across the whole network of sending a message from each node.

In order to calculate the lifetime remaining due to a routing scheme we require additional intrinsic information about the nodes, namely the charge Q_i remaining in the battery and the quiescent energy consumption per reporting cycle E_i due to constant micro-controller operation, sensor measurements, running an on-board display, etc. The life of the current node therefore is modelled as

$$L_i = \frac{Q_i}{(E_i + C_i)N} \quad (4)$$

where N is the number of reporting cycles per unit time. We emphasise $L_i = L_i(R)$, i.e. L_i is a function of all the routes which utilise v_i .

Our goal is to prolong the average life of the network, that is to minimise the total energy consumed, and to maximise the time before any individual node requires its battery to be recharged or changed. We therefore seek to maximise the two objective problem:

$$\text{Maximise } f_1(R) = \frac{1}{n} \sum_{i=1}^n L_i(R). \quad (5)$$

$$\text{Maximise } f_2(R) = \min_{i \in [1, n]} L_i(R). \quad (6)$$

In addition, it may be important to maximise the lifetime of one or more nodes v_i for $i \in U$, because, for example, they are particularly inaccessible. In this case the two-objective problem is augmented with a third objective:

$$\text{Maximise } f_3(R) = \min_{i \in U} L_i(R). \quad (7)$$

Solving this multi-objective problem may result in multiple solutions, as opposed to a single solution for single objective optimisation. In this case, there exists a set of solutions which are Pareto optimal; that is, there are no other feasible solutions available that improve performance on one objective, without a simultaneous decrease in at least in one other objective [3].

The dominance criterion is used to locate such solutions in the search space. The dominance criterion from a routing optimisation perspective is described as follows. In a multi-objective problem with D objectives, a routing scheme, R' ,

is said to dominate another routing scheme, R , denoted $R' \succ R$, iff

$$\begin{aligned} f_i(R') &\geq f_i(R) \quad \forall i = 1, 2, \dots, D, \text{ and} \\ f_i(R') &> f_i(R) \text{ for some } i. \end{aligned} \quad (8)$$

Hence, we seek the maximal set of feasible routes which are mutually non-dominating, which is known as the Pareto set, \mathcal{P} .

3. SEARCH SPACE PRUNING

The multi-objective optimisation problem described above is a combinatorial optimisation problem with, for practical WSNs, a vast number of potential solutions. It is crucial for practical implementations that the optimisation process is fast. A way to improve the speed of optimisation is to sensibly prune the search space, while retaining important potential solutions. For this purpose we deploy a novel k -shortest path based technique. Before describing the evolutionary algorithm employed, we therefore discuss the way in which solutions are represented and the search space pruned to permit an efficient approximation of the Pareto set.

The number of possible routing schemes, i.e. the search space size, depends on the number of available routes for each node in the system. For instance, in a network with n nodes excluding the base station with let the number of available loopless paths from v_i to v_B be a_i . In this case, the number of possible routing schemes, i.e. the number of combinations of routes for individual nodes that can build the routing scheme, is:

$$Z_a = \prod_{i=1}^n a_i. \quad (9)$$

In order to combat the potential growth in the size of the search space as the number of nodes increases we therefore limit the number of potential routes available to each node. More specifically, we attempt to search for solutions in the space defined by the k -shortest paths for each node, where the metric defining the k -shortest paths is defined below. This reflects our intuition that short paths to the base station are most likely to be energy efficient. We select from among several shortest path routes for each node because if each node were to utilise its shortest path a single node on many of them would be disproportionately burdened.

Algorithms for discovering the shortest path between two nodes in a weighted graph are well known and the shortest path can be found in $O(n \log n)$ time. However, as we noted above, the energy costs in this problem are associated with the nodes themselves rather than with the edges. We therefore weight the edges in the network graph to associate the energetic cost at the nodes with the edges connecting them. Consider the nodes v_i and v_j . We define the weight of the edge between them as:

$$w_{ij} = \frac{e_{ij}}{Q_i} + \frac{e_{ji}}{Q_j}, \quad (10)$$

where, as above, e_{ij} is the energy required to transmit a message from v_i to v_j , and Q_i and Q_j are the battery charges. It is expected that $e_{ij} = e_{ji}$. This edge weighting models the fact that a high transmission cost can be borne by nodes with a high battery charge, but transmission is relatively expensive for nodes with low battery charge because each transmission will make a larger fractional depletion of the




Figure 2: Correlation between objectives and composite cost. Scatter plots show (*right*) the composite cost and true average lifetime $f_1(R)$ and (*left*) minimum lifetime $f_2(R)$ for 1000 randomly generated routing schemes with random residual charges at the nodes.

charge. Likewise transmissions are free if $Q_i = \infty$ (i.e. the node is connected to mains power). Note that $Q_i / \sum_j e_{ij}$ is an estimate of the lifetime of v_i . We call the cost of a routing scheme calculated using the weights w_{ij} the *composite cost*.

As we require diversity in the search space and the possibility of load balancing among nodes, we propose to evolve solutions from among the k -shortest paths for each node calculated with the composite cost (10). A number of algorithms are available for computing the k shortest paths; see for example [12, 5]. In our implementation we have used Eppstein’s algorithm [5] modified to produce only simple or loopless paths, which matches the best upper-bound time complexity for finding k -shortest simple paths [5]. We denote the m th shortest route found for node v_i by S_i^m , for $m = 1, \dots, k$. Using this technique, the total size of the search space is no larger than k^n , although it may be smaller than this because nodes close to v_B may not have k loopless paths to v_B .

3.1 Relationship between objectives and composite cost

In order to validate the use of the composite metric for choosing candidate routes, Figure 2 presents the correlation between the objectives and composite cost for 1000 randomly chosen routing schemes with random residual charge levels Q_i at nodes for the Victoria & Albert museum network described below. There is a fairly strong negative correlation between the average battery life $f_1(R)$ and the composite cost. This indicates that routing schemes selected with low composite cost via the k -shortest paths algorithm are likely to have long average lives. On the other hand, there exists a weak negative relationship between the minimum lifetime and composite cost. This is unsurprising since the minimum lifetime depends on a *single* node in the network and is highly dependent on the interaction between routes.

Algorithm 1 Optimisation of Battery Powered Mesh Network

Inputs

- 1: T = Number of iterations
- 2: s = Size of initial archive
- 3: μ = Perturbation rate
- 4: c = Uniform crossover rate

Steps

- 1: $A \leftarrow \text{InitialiseArchive}(s)$ \triangleright Initialise random archive
 - 2: **for** $i = 1 \rightarrow T$ **do**
 - 3: $\{R_1, R_2\} \leftarrow \text{Select}(A)$ \triangleright Select two parent solutions
 - 4: $R' \leftarrow \text{UniformCrossOver}(R_1, R_2, c)$
 - 5: $R'' \leftarrow \text{Perturb}(R', \mu)$ \triangleright Mutation
 - 6: $A \leftarrow \text{NonDominated}(A \cup R'')$ \triangleright Update archive
 - 7: **end for**
 - 8: **return** A \triangleright Approximation of the Pareto set
-

4. MULTI-OBJECTIVE ROUTING OPTIMISATION

The multi-objective evolutionary algorithm used in this approach is a real-valued genetic algorithm, which maintains an unconstrained Pareto archive to reap the benefits of better convergence properties [6], but does not employ an independent search population. Algorithm 1 describes this process in more detail.

Solutions R are represented by vectors of n integers; the i th element of the solution indexes one of the k shortest paths found for node v_i . In the initialisation step, we generate at random a population from the pruned search space. This population comprises random routing schemes where the member routes are picked from the k -shortest paths for each node; thus for node v_i the shortest paths are selected from $\{S_i^m\}_{m=1}^k$. Additionally, we include in the initial population the routing scheme which uses the first shortest route for each node. The initial archive of non-dominated solutions A is the maximal non-dominated subset of this random population. At any step of the evolution, A is the current approximation of the Pareto set.

During the evolution process, we randomly select two routing schemes. These parent routing schemes are then crossed-over uniformly with a crossover rate c , resulting in the single dominating offspring. The child is then perturbed by changing routes randomly within the routing scheme, where the number of changes depends on the perturbation rate, μ . The offspring from perturbation is compared against the members in the archive: if it is not dominated by any of them then it enters the archive and any elements of A which are dominated by the new solution are removed from A . In this fashion only the non-dominated routing schemes are preserved in the archive and the archive can only approach the true Pareto set. The process of evolution continues for a fixed number of episodes (alternatively, another termination criterion, such as specified minimum dominated hypervolume, may be employed).

Once the evolution process is finished, the decision maker may manually choose the operating point using the final approximation of the Pareto front, based on the expected network longevity. The chosen routing scheme is sent to the nodes through the base station, and it then becomes the active data reporting scheme in the system.

5. ILLUSTRATION

A real network deployed in the Victoria & Albert Museum, London, is used to illustrate the proposed approach. In a controlled environment over a vast area, such as Victoria & Albert Museum, it is essential to monitor temperature and humidity in galleries and display cases for the preservation of the artefacts. Compared to wired networks, battery powered WSNs carry huge advantages in deployment cost and flexibility.

The network incorporated 30 sensor nodes and a base station, spanning five floors within an approximate area of 35000 m^2 . This provides a challenging radio environment with thick, solid walls and a dynamic medium varying with the rate of passage of people in the galleries.

The connectivity map was built in an initial mapping phase in which nodes pinged each other using a range of baud rates and powers to discover the minimum energy configuration for communication between those nodes within radio range of each other.

Subspace pruning using $k = 10$ shortest paths still results in a search space of 10^{30} solutions. The initial population was built with 100 randomly chosen solutions from this subspace together with the solution consisting of the shortest route for each node. The initial archive was then found by extracting the non-dominated solutions in this population.

For optimisation purposes, we used $\mu = c = 0.1$ as the perturbation and crossover rates; these rates were chosen after a short empirical study on simulated networks, but the performance of the optimiser is relatively insensitive to their precise values. Using the dominated hypervolume measure, the Pareto set was well converged after 150,000 iterations.

A single core Python implementation takes about 2 minutes to complete 150,000 iterations on a 2.5 GHz machine. As a consequence, the optimum routes for the new installations can be found readily, making this approach particularly feasible as part of a real system.

5.1 Baseline optimisation

Figure 3 shows the Pareto front resulting from the optimisation of the network. As the figure shows, routing schemes in the Pareto set provide the network operator with a wide range of routes trading off the average lifetime of the network against the time before any single node needs its battery replacing. We note that the optimisation has resulted in routing schemes that are substantially better, in terms of both average and minimum lifetime, than routing schemes utilising randomly chosen routes from the 10 shortest paths (I , marked with blue crosses in the figure). Indeed, the routing scheme with the longest minimum life has an average lifetime that is better than any of the solutions in the initial random population. Interestingly, the routing scheme, R_c in which each node uses its shortest path route (according to the composite cost) lies close to the Pareto front, although it is dominated by solutions in the Pareto front. We include R_c in the initial archive.

As a result of optimisation, we were able to derive the trade-off front that dominates the shortest paths routing scheme, and thus any solution from the front performs better in comparison. For example, the solution directly above R_c can provide 1 more month in the minimum lifetime, with the same average lifetime. Also, considering the left-most solution, we can get an improvement of 2.5 months in minimum lifetime at the expense of 0.5 months of average lifetime.




Figure 3: Baseline Optimisation. Pareto front approximation (red) and randomly generated initial solutions (blue) for the Victoria & Albert museum network. Non-dominated solutions comprising the initial archive are shown as in blue with green rings and the solution in which each node uses its shortest path based on the composite cost is shown in orange/black.

6. EXTENDING MINIMUM LIFETIME

It is of considerable practical importance to maximise the minimum lifetime of any node in the network, because the network requires no maintenance until the battery of the shortest-lived node is exhausted. Rather than using a single routing scheme for the lifetime of the network it can be advantageous to adopt a new routing scheme partway through its life which can serve to prolong the time before any single node's battery is exhausted. Rerouting of the network can be accomplished at low energy cost from the base station. To see how this can be useful, suppose that under the initial routing scheme R , a node v^* has the minimum lifetime. Now suppose that the network is optimised a second time after it has been in operation for a while. The result of this second optimisation can be a routing scheme in which v^* is very lightly loaded, prolonging its life, while a different node, v' , which was lightly loaded in the first epoch, now carries a heavier load. In this way the life of v^* is extended and it may be that v^* or v' or some other node is exhausted first, but in any case the time before any single node is exhausted may be extended.

To illustrate this, we choose the routing solution with the best minimum lifetime from the Pareto front shown in Figure 3, that is, the left-most solution, which we denote by R . We simulate the operation of the network for 6 months, roughly one third of the average lifetime of the network or roughly half of the time (1.12 years) before $v^* = v_{21}$ would be exhausted. Note that during this time, the batteries in the various nodes are discharged at different rates, so that after 6 months the residual charges $\{Q_i\}_{i=1}^n$ are unequal. Figure 4 shows the result of re-optimising the routings using these residual charges. Here A' is the new Pareto optimal front, which can be seen to dominate the original routing scheme R which is marked in blue. Note that the remaining lifetimes for R are 6 months smaller than those appearing in Figure 3 because the network has been running for 6 months.

Minimum lifetime: 0.62 years for node 21

Minimum lifetime: 0.67 years for node 21

Minimum lifetime: 0.62 years for node 29

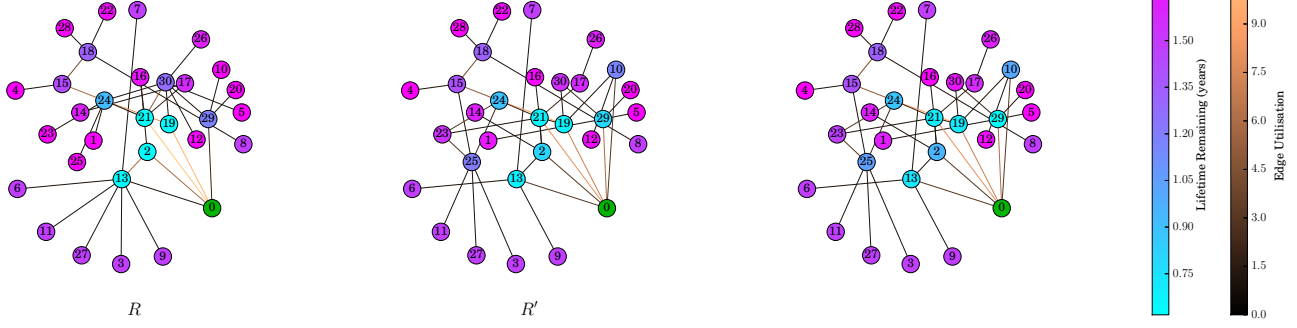


Figure 6: Network graphs. Nodes are coloured according to the remaining battery life and edges are coloured according to the utilisation of the corresponding link. *Left*: Network corresponding to the minimum lifetime solution R shown in the Figure 3. *Middle*: Network R' , the best minimum life network following re-optimisation based on R after 6 months. *Right*: Routing R' showing lifetime remaining after 6 months operation, but using R' throughout.




Figure 4: Re-optimised Pareto Front. The system simulated for 6 months with operating point R (blue cross). Re-optimisation with battery states after 6 months of operating R generates the Pareto front A' (red dots), with R' being the best minimum lifetime solution.

We draw attention to the best minimum lifetime solution, R' from this second optimisation which has a significantly longer life than the scheme R . Thus using R for the first 6 months, followed by R' prolongs the life of the network before any battery needs changing.

To examine this further, Figure 5 compares the overall lifetimes available using the two-stage optimisation with the original lifetimes. The original Pareto front is shown in red, while the green symbols show the Pareto archive A' resulting from the second optimisation (Figure 4) with 6 months added to each objective coordinate to show the total lifetimes of the two-stage networks. Also shown, in blue and labelled \hat{A}' , are the routing schemes in A' re-evaluated using the batteries with charges corresponding to the initial optimisation. The \hat{A}' routing schemes do not dominate the original schemes in A with one or two minor exceptions which indicate that the Pareto front is not quite fully converged.






Figure 5: Comparison of overall lifetimes from one and two-stage optimisation. The Pareto front corresponding to the original optimisation is shown in red, with R being the best minimum lifetime solution from this archive A . Re-optimisation with battery states after 6 months of operating R generates the trade-off of overall lifetimes A' (green pluses), with R' being the best minimum lifetime solution. Blue crosses indicate the lifetimes of solutions from A' evaluated with the initial battery states; \hat{R}' is the solution R' evaluated with the initial battery states.

We note that the re-optimisation was performed using k -shortest paths derived using battery charges after 6 months operation of R so a slightly different set of potential solutions was available to this optimisation.

In Figure 6 we portray network graphs corresponding to R , R' , and \hat{R}' . The left panel of the figure illustrates the lifetimes of the nodes using the best minimum lifetime routing R from the original optimisation after it has been in use for 6 months. Nodes are coloured according to their remaining life and in this case it can be seen that node 21 has the minimum lifetime of 0.62 years. Edges are shaded according to their utilisation and, unsurprisingly, it can be seen the



(a) Re-optimisation following deletion of edge $v_{19} \rightarrow v_B$.



(b) Re-optimisation following failure of node v_{19} .

Figure 7: Dynamic Optimisation. Average and minimum lifetimes remaining for solutions in the Pareto front after 6 months of operation (magenta pluses) using the route scheme R_d (black/orange diamond). The Pareto front resulting from re-optimisation as a result of node or edge failure is shown by red dots. The routes repaired by random selection from k shortest paths are denoted by blue crosses, and initial archive by green rings.

links to v_{21} are well used (the link $v_{21} \rightarrow v_B$ is used by 11 routes), resulting in its short life.

The network illustrated in the middle panel of Figure 6 corresponds to R' , the routing scheme with the best minimum lifetime resulting from a second optimisation after 6 months based on R . As can be seen, although v_{21} is still the first node to be exhausted its life has been extended by heavier utilisation of nodes 29 and 10; in particular the $v_{21} \rightarrow v_B$ link is now used only 8 times. This redistribution of loading during the two epochs is what allows the combination of R and R' to provide superior performance to R alone. The right panel of Figure 6 shows the lifetimes remaining after 6 months using R' from the outset. Here v_{29} has the minimum life of 0.62 years which is equivalent to the minimum life v_{21} using R and no benefit is accrued from using either scheme alone. It is the combination of the two schemes that allows energy to be expended more evenly around the network, protecting different nodes in different epochs, that enables the overall minimum life of the network to be prolonged.

Clearly, this two-stage combination of routing schemes could be extended to incorporate additional epochs, although we speculate that smaller relative gains will be made as additional epochs are added. Also in the work we present here the point at which to switch from R to R' was fixed in advance of the optimisations and the second optimisation was carried out independently of the first. Current work is investigating simultaneous optimisation of two or more epoch combinations.

7. DYNAMIC OPTIMISATION

As WSNs operate in a dynamic environment, it is conceivable that over time some of the links in the routing scheme will fail. In addition nodes can disappear from the network due to hardware failure or changing requirements and new nodes can be introduced as the network is extended. Hence,

it is important to detect failures and re-establish connectivity with optimal routes, thus enabling self-healing within the network.

These failures can be detected from the failure of nodes to report data within an expected period of time as indicated by the packet loss ratio. The particular edge that failed can be detected as nodes succeeding the failed link send a distress signal to the base station. In such cases, the solutions from the original Pareto front may become infeasible, as they may include the lost nodes or edges. Therefore, depending on the reliability thresholds as set by the decision maker, a re-optimisation is necessary in order to find the optimal front in this new state using the latest connectivity information. We note that the solutions from the original Pareto set contain previous knowledge of good routes which is useful for routes unaffected by the failure. These routes can be fused with newly-generated random routes in order to build the initial archive for a re-optimisation. Here we generate an initial archive for re-optimisation by replacing the affected routes from the original front with randomly selected alternatives from the available $k = 10$ shortest paths and extracting the non-dominated subset. Using this initial archive can promote subsequent rapid optimisation.

To illustrate dynamic optimisation, we select a solution R_d , which provides a good balance between minimum lifetime and average lifetime among the solutions in the original Pareto front. The network is then simulated for 6 months, and consequently R_d has a minimum lifetime of 0.47 years and an average lifetime of 1.47 years as shown in Figure 7. We noted that the node v_{19} has the minimum lifetime at this stage.

To exemplify the effects of edge failure, we delete the link between v_{19} and the base station. This edge was selected as it is common among the solutions in the original front with an average of 13 routes per solution using this link; the failure of this link therefore represents considerable damage to the network. As a result, all solutions in the original

Pareto set become infeasible. By combining the unaffected routes with random routes from 10 shortest paths, we derive an initial archive, I_e for a new optimisation. This archive is then used to optimise the network as described in section 5.1.

The Pareto fronts resulting from the re-optimisation after damage are shown in Figure 7. For both edge deletion (Figure 7a) and node failure (Figure 7b), a large portion of the new optimal front is dominated by the original front. This is unsurprising because in both cases the solution space has been significantly reduced and new optimal routes to the base station are necessarily longer than the original routes. Interestingly, however, the re-optimisation has found solutions that prolong the minimum life of the network in the way as discussed in section 6. This leads to a considerable portion of the re-estimated front being mutually non-dominating with solutions from the original Pareto front, giving the network manager considerable flexibility to choose a new operating point. We remark that because both objectives are expressed in common units (years) it is straightforward to automatically choose the new routing scheme that is closest to the original R_d using the Euclidean distance.

The utility of initialising the re-optimisation archive with solutions constructed from the original archive is evident in that, based on the hypervolume, the re-optimisation converged roughly three times as fast as an optimisation starting from random solutions, even in this case where the network was considerably damaged.

8. CONCLUSIONS

Using mesh network topologies in low power Wireless Sensor Networks poses a challenging problem of finding routes that best preserve the lifetime of individual nodes and the network as a whole. In this paper we have proposed a multi-objective evolutionary approach to approximate the optimum trade-off between local and global objectives, i.e. minimum lifetime remaining for any node and average lifetime remaining for the whole network. The potentially vast solution space is reduced by searching for routes among the approximate k -shortest paths from each node to the base station. This is made possible by associating energy costs at nodes with the edges between them using a heuristic metric. Use of the approximate k -shortest paths also provides a convenient method of representing solutions in the evolutionary optimiser.

We have shown in a real system, deployed in the Victoria & Albert Museum, that the optimiser is able to estimate the optimum trade-off between the minimum lifetime of any node in the network and the average lifetime of the network as a whole. Interestingly it was also found that the range of minimum lifetimes available can also be significantly extended by using more than one routing scheme successively. In essence this works by protecting during the first epoch nodes that will bear a heavy load in the second epoch and *vice versa*. To achieve these combined routing schemes two optimisations were required and a solution from the first optimisation must be selected as the basis for the second optimisation. Current work focuses on efficient evolutionary optimisation methods to simultaneously locate solutions for both epochs which has the potential to achieve more efficient networks.

9. ACKNOWLEDGEMENTS

This work was supported by a Knowledge Transfer Partnership awarded to the University of Exeter and The IMC Group Ltd. We would like to thank Martin Hancock (Technical Director, The IMC Group Ltd.), and Neil Lundy (Engineering Manager, The IMC Group Ltd.) for their support and contributions. We would also like to thank Boris Pretzel (Chief Scientist, Victoria & Albert Museum), and Bhavesh Shah (Scientist, Victoria & Albert Museum), for facilitating the real network test. A patent application based on this work has been made (reference number: GB 1400353.7).

10. REFERENCES

- [1] S. Bashyal and G. K. Venayagamoorthy. Collaborative routing algorithm for wireless sensor network longevity. In *Intelligent Sensors, Sensor Networks and Information, ISSNIP'07.*, pages 515–520, 2007.
- [2] A. W. Brander and M. C. Sinclair. A comparative study of k -shortest path algorithms. In *Proc. 11th UK Performance Eng. Workshop*, pages 370–379, 1995.
- [3] C. A. C. Coello, G. B. Lamont, and D. A. V. Veldhuisen. *Evolutionary algorithms for solving multi-objective problems*. Springer, 2001.
- [4] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2001.
- [5] D. Eppstein. Finding the k shortest paths. *SIAM Journal on Computing*, 28(2):652–673, 1998.
- [6] J. E. Fieldsend, R. M. Everson, and S. Singh. Using unconstrained elite archives for multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 7(3):305–323, 2003.
- [7] A. Förster. Machine learning techniques applied to wireless ad-hoc networks: Guide and survey. In *Intelligent Sensors, Sensor Networks and Information, ISSNIP'07*, pages 365–370, 2007.
- [8] O. Islam and S. Hussain. An intelligent multi-hop routing for wireless sensor networks. In *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology Workshops*, pages 239–242, 2006.
- [9] S. Kamath and A. Nasipuri. Integrated load balanced and energy aware routing in large scale wireless sensor networks. In *IEEE, 30th International Performance Computing and Communications Conference (IPCCC), 2011*, pages 1–8, 2011.
- [10] R. V. Kulkarni, A. Förster, and G. K. Venayagamoorthy. Computational intelligence in wireless sensor networks: A survey. *IEEE Comms Surveys & Tutorials*, 13(1):68–96, 2011.
- [11] F. Xue, A. Sanderson, and R. Graves. Multi-objective routing in wireless sensor networks with a differential evolution algorithm. In *Proceedings of the 2006 IEEE International Conference on Networking, Sensing and Control, ICNSC '06*, pages 880–885, 2006.
- [12] J. Y. Yen. Finding the K shortest loopless paths in a network. *Management Science*, 17:712–716, 1971.
- [13] H. Yetgin, K. T. K. Cheung, and L. Hanzo. Multi-objective routing optimization using evolutionary algorithms. In *Wireless Communications and Networking Conference (WCNC), pages 3030–3034, 2012*.